

The Status of the P versus NP Problem

Lance Fortnow
Northwestern University

1. INTRODUCTION

When Moshe Vardi asked me to write this piece for CACM, my first reaction was the article could be written in two words

Still open.

When I started graduate school in the mid-1980's, many believed that the quickly developing area of circuit complexity would soon settle the **P** versus **NP** problem, whether every algorithmic problem with efficiently verifiable solutions have efficiently computable solutions. But circuit complexity and other approaches to the problem have stalled and we have little reason to believe we will see a proof separating **P** from **NP** in the near future.

Nevertheless the computer science landscape has dramatically changed in the nearly four decades since Steve Cook presented his seminal **NP**-completeness paper *The Complexity of Theorem-Proving Procedures* [11] in Shaker Heights, Ohio in early May, 1971. Computational power has dramatically increased, the cost of computing has dramatically decreased, not to mention the power of the Internet. Computation has become a standard tool in just about every academic field. Whole subfields of biology, chemistry, physics, economics and others are devoted to large-scale computational modeling, simulations and problem solving.

As we solve larger and more complex problems with greater computational power and cleverer algorithms, the problems we cannot tackle begin to stand out. The theory of **NP**-completeness helps us understand these limitations and the **P** versus **NP** problems begins to loom large not just as an interesting theoretical question in computer science, but as a basic principle that permeates all the sciences.

So while we don't expect the **P** versus **NP** problem to be resolved in the near future, the question has driven research

in a number of topics to help us understand, handle and even take advantage of the hardness of various computational problems.

In this survey we will look at how people have tried to solve the **P** versus **NP** problem but also how this question has shaped so much of the research in computer science and beyond. We will look at how to handle **NP**-complete problems and the theory that has developed from those approaches. We show how a new type of "interactive proof systems" led to limitations of approximation algorithms. We consider whether quantum computer can solve **NP**-complete problems (short answer: not likely). We end describing a new long-term project that will try to separate **P** from **NP** using algebraic-geometric techniques.

This survey won't try to be 100% technically or historically accurate or complete but rather informally describe the **P** versus **NP** problem and the major directions in computer science inspired by this question over the past several decades.

2. WHAT IS THE P VERSUS NP PROBLEM?

Suppose we have a large group of students that we need to pair up to work on projects. We know which students are compatible with each other and we want to put them in compatible groups of two. We could search all possible pairings but even for 40 students we would have more than three hundred billion trillion possible pairings.

In 1965, Jack Edmonds [13] gave an efficient algorithm to solve this matching problem and suggested a formal definition of "efficient computation" (runs in time a fixed polynomial of the input size). The class of problems with efficient solutions would later become known as **P** for "Polynomial Time".

But many related problems do not seem to have such an efficient algorithm. What if we wanted to make groups of three students with each pair of students in each group compatible (Partition Into Triangles)? What if we wanted to find a large group of students all of whom are compatible with each other (Clique)? What if we wanted to sit the students around a large round table with no incompatible students sitting next to each other (Hamiltonian Cycle)? What if we put the students into three groups so that each student is in the same group with only his or her compatibles (3-Coloring)?

All these problems have a similar flavor: Given a potential solution, for example a seating chart for the round table, we can validate that solution efficiently. The collection of problems that have efficiently verifiable solutions is known as **NP** (for “Nondeterministic Polynomial-Time” if you have to ask).

So $\mathbf{P} = \mathbf{NP}$ means that for every problem that has an efficiently verifiable solution, we can find that solution efficiently as well.

We call the very hardest **NP** problems (which include Partition Into Triangles, Clique, Hamiltonian Cycle and 3-Coloring) “**NP**-complete”, i.e. given an efficient algorithm for one of them, we can find an efficient algorithm for all of them and in fact any problem in **NP**. Steve Cook, Leonid Levin and Richard Karp [11, 28, 25] developed the initial theory of NP-completeness that generated multiple ACM Turing Awards.

In the 1970’s, theoretical computer scientists showed hundreds more problems **NP**-complete (see [17]). An efficient solution to any **NP**-complete problem would imply $\mathbf{P} = \mathbf{NP}$ and an efficient solution to every **NP**-complete problem.

Most computer scientists quickly came to believe $\mathbf{P} \neq \mathbf{NP}$ and trying to prove it quickly became the single most important question in all of theoretical computer science and one of the most important in all of mathematics. Soon the **P** versus **NP** problem became an important computationally issue in nearly every scientific discipline.

As computers grew cheaper and more powerful, computation started playing a major role in nearly every academic field, especially the sciences. The more scientists can do with computers, the more they realize some problems seem computationally difficult. Many of these fundamental problems turn out to be NP-complete. A small sample:

- Finding a DNA sequence that best fits a collection of fragments of the sequence (see [21]).
- Finding a ground state in the Ising model of phase transitions (see [9]).
- Finding Nash Equilibriums with specific properties in a number of environments (see [10]).
- Finding optimal protein threading procedures [27].
- Determining if a mathematical statement has a short proof (follows from Cook [11]).

In 2000, the Clay Math Institute named the **P** versus **NP** problems as one of the seven most important open questions in mathematics and has offered a million-dollar prize for a proof that determines whether or not $\mathbf{P} = \mathbf{NP}$.

3. WHAT IF $\mathbf{P} = \mathbf{NP}$?

To understand the importance of the **P** versus **NP** problem let us imagine a world where $\mathbf{P} = \mathbf{NP}$. Technically we could have $\mathbf{P} = \mathbf{NP}$ but not have practical algorithms for most **NP**-complete problems. But suppose in fact that we do have very quick algorithms for all these problems.

Many focus on the negative, that if $\mathbf{P} = \mathbf{NP}$ then public-key cryptography becomes impossible. True but what we will gain from $\mathbf{P} = \mathbf{NP}$ will make the whole internet look like a footnote in history.

Since all the **NP**-complete optimization problems become easy, everything will be much more efficient. Transportation of all forms will be scheduled optimally to move people and goods around quicker and cheaper. Manufacturers can improve their production to increase speed and create less waste. And I’m just scratching the surface.

Learning becomes easy by using the principle of Occam’s razor—we simply find the smallest program consistent with the data. Near perfect vision recognition, language comprehension and translation and all other learning tasks become trivial. We will also have much better predictions of weather and earthquakes and other natural phenomenon.

$\mathbf{P} = \mathbf{NP}$ would also have big implications in mathematics. One could find short fully logical proofs for theorems but these fully logical proofs are usually extremely long. But we can use the Occam razor principle to recognize and verify mathematical proofs as typically written in journals. We can then find proofs of theorems that have reasonably length proofs say in under 100 pages. A person who proves $\mathbf{P} = \mathbf{NP}$ would walk home from the Clay Institute not with one million-dollar check but with seven (actually six since the Poincaré Conjecture appears solved).

Don’t get your hopes up. Complexity theorists generally believe $\mathbf{P} \neq \mathbf{NP}$ and such a beautiful world cannot exist.

4. APPROACHES TO SHOWING $\mathbf{P} \neq \mathbf{NP}$

In this section we present a number of ways we have tried and failed to prove $\mathbf{P} \neq \mathbf{NP}$. The survey of Fortnow and Homer [15] gives a fuller historical overview of these techniques.

4.1 Diagonalization

Can we just construct an **NP** language L specifically designed so that every single polynomial-time algorithm fails to compute L properly on some input? This approach, known as diagonalization, goes back to the 19th Century.

In 1874 Cantor [8] showed that the real numbers are uncountable using a technique known as diagonalization. Given a countable list of reals, Cantor showed how to create a new real number not on that list.

Turing, in his seminal paper on computation [39], used a similar technique to show that the Halting problem is not computable. In the 1960’s complexity theorists used diagonalization to show that given more time or memory one can solve more problems. Why not use diagonalization to separate **NP** from **P**?

Diagonalization requires simulation and we don’t know how a fixed **NP** machine can simulate an arbitrary **P** machine. Also a diagonalization proof would likely relativize, i.e., work even if all machines involved have access to the same additional information. Baker, Gill and Solovay [7] showed no

relativizable proof can settle the \mathbf{P} versus \mathbf{NP} problem in either direction.

Complexity theorists have used diagonalization techniques to show some \mathbf{NP} -complete problems like Boolean formula satisfiability cannot have algorithms that use both a small amount of time and memory [40] but this is a long way from $\mathbf{P} \neq \mathbf{NP}$.

4.2 Circuit Complexity

To show $\mathbf{P} \neq \mathbf{NP}$ it is sufficient to show some \mathbf{NP} -complete problem cannot be solved by relatively small circuits of AND, OR and NOT gates (the number of gates bounded by a fixed polynomial in the input size).

In 1984, Furst, Saxe and Sipser [16] showed that small circuits cannot solve the parity function if the circuits have a fixed number of layers of gates. In 1985, Razborov [32] showed the \mathbf{NP} -complete problem of finding a large clique does not have small circuits if one only allows AND and OR gates (no NOT gates). If one extends Razborov's result to general circuits one will have proved $\mathbf{P} \neq \mathbf{NP}$.

Razborov later showed his techniques would fail miserably if one allows NOT gates [33]. Razborov and Rudich [34] develop a notion of "natural" proofs and give evidence that our limited techniques in circuit complexity cannot be pushed much further. And in fact we haven't seen any significantly new circuit lower bounds in the past twenty years.

4.3 Proof Complexity

Consider the set of Tautologies, the Boolean formulas ϕ of variables over ANDs, ORs and NOTs such that every setting of the variables to True and False makes ϕ true, for example the formula

$$(x \text{ AND } y) \text{ OR } (\text{NOT } x) \text{ OR } (\text{NOT } y).$$

A literal is a variable or its negation, i.e. x or $\text{NOT } x$. A formula, like the one above, is in Disjunctive Normal Form (DNF) if it is the OR of ANDs of one or more literals.

If a formula ϕ is not a tautology, we can give an easy proof of that fact by exhibiting an assignment of the variables that makes ϕ false. But if ϕ were indeed a tautology we don't expect short proofs of that. If one could prove there are no short proofs of tautology that would imply $\mathbf{P} \neq \mathbf{NP}$.

Resolution is a standard approach to proving tautologies of DNFs by finding two clauses of the form $(\psi_1 \text{ AND } x)$ and $(\psi_2 \text{ AND NOT } x)$ and adding the clause $(\psi_1 \text{ AND } \psi_2)$. A formula is a tautology exactly when one can produce an empty clause in this manner.

In 1985, Haken [22] showed that tautologies that encode the pigeonhole principle ($n + 1$ pigeons in n holes means some hole has more than one pigeon) do not have short resolution proofs.

Since then complexity theorists have shown similar weaknesses in a number of other proof systems including cutting planes, algebraic proof systems based on polynomials and restricted versions of proofs using the Frege axioms, the basic axioms one learns in an introductory logic course.

But to prove $\mathbf{P} \neq \mathbf{NP}$ we would need to show that tautologies cannot have short proofs in an arbitrary proof system. Even a breakthrough result showing tautologies don't have short general Frege proofs would not suffice in separating \mathbf{NP} from \mathbf{P} .

5. DEALING WITH HARDNESS

So you have an \mathbf{NP} -complete problem you just have to solve. If as we believe $\mathbf{P} \neq \mathbf{NP}$ you won't find a general algorithm that will correctly and accurately solve your problem all of the time. But sometimes you need to solve the problem anyway. All hope is not lost. In this section we describe some of the tools one can use on \mathbf{NP} -complete problems and how computational complexity theory studies these approaches. Typically one needs to combine several of these approaches when tackling \mathbf{NP} -complete problems in the real world.

5.1 Brute Force

Computers have gotten faster, much much faster since \mathbf{NP} -completeness was first developed. Brute force search through all possibilities is now possible for some small problem instances. With some clever algorithms we can even solve some moderate size problems with ease.

The \mathbf{NP} -complete traveling salesperson problem asks for the smallest distance tour through a set of specified cities. Using extensions of the cutting-plane method we can now solve, in practice, traveling salespeople problems with more than ten thousand cities (see [4]).

Consider the 3SAT problem, solving Boolean formula satisfiability where formulas are in the form of the AND of several clauses where each clause is the OR of three literal (variables or negations of variables). 3SAT remains \mathbf{NP} -complete but the best algorithms can in practice solve SAT problems on about 100 variables. We have similar results for other variations of satisfiability and many other \mathbf{NP} -complete problems.

But for satisfiability on general formulae and on many other \mathbf{NP} -complete problems we do not know algorithms better than essentially searching all the possibilities. In addition all these algorithms have exponential growth in their running times, so even a small increase in the problem size can kill what was an efficient algorithm. Brute force alone will not solve \mathbf{NP} -complete problems no matter how clever we are.

5.2 Parameterized Complexity

Consider the Vertex Cover problem, find a set of k "central people" such that for every compatible pair of people, at least one of them is central. For small k we can determine whether a central set of people exists efficiently no matter the total number n of people we are considering. For the Clique problem even for small k the problem can still be difficult.

Downey and Fellows [12] developed a theory of parameterized complexity that gives a fine-grained analysis of the complexity of \mathbf{NP} -complete problems based on their parameter size.

5.3 Approximation

We cannot hope to solve **NP**-complete optimization problems exactly but often we can get a good approximate answer. Consider the traveling salesperson problem again with distances between cities given as the crow flies (Euclidean distance). This problem remains **NP**-complete but Arora [5] gives an efficient algorithm that gets very close to the best possible route.

Consider the MAX-CUT problem of dividing people into two groups to maximize the number of incompatibles between the groups. Goemans and Williamson [18] uses semi-definite programming to give a division of people only a .878567 factor of the best possible.

In Section 6 we will see how a theory of probabilistic proof systems leads to provable limitations of approximation algorithms.

5.4 Heuristics and Average-Case Complexity

The study of **NP**-completeness focuses on how algorithms perform on the worst possible inputs. However the specific problems that arise in practice may be much easier to solve. Many computer scientists employ various heuristics to solve **NP**-complete problems that arise from the specific problems in their fields.

While we create heuristics for many of the **NP**-complete problems, Boolean formula Satisfiability (SAT) receives more attention than any other. Boolean formulas, especially those in conjunctive normal form (CNF), the AND of ORs of variables and their negations, have a very simple description and yet are general enough to apply to a large number of practical scenarios particularly in software verification and artificial intelligence. Most natural **NP**-complete problems have simple efficient reductions to the satisfiability of Boolean formula. In competition these SAT solvers can often settle satisfiability of formulas of a million variables [1].

Computational complexity theorists study heuristics by considering average-case complexity—how well can algorithms perform on average from instances generated by some specific distribution.

Levin [29] developed a theory of efficient algorithms over a specific distribution and formulated a distributional version of the **P** versus **NP** problem.

Some problems like versions of the shortest vector problem in a lattice or computing the permanent of a matrix are hard on average exactly when they are hard on worst-case inputs, but neither of these problems is believed to be **NP**-complete. Whether similar worst-to-average reductions hold for **NP**-complete sets is an important open problem.

Average-case complexity plays an important role in many areas of computer science, particularly cryptography which we discuss in Section 7.1.

6. INTERACTIVE PROOFS AND LIMITS OF APPROXIMATION

In Section 5.3 we saw how sometimes one can get good approximate solutions to **NP**-complete optimization problems.

Many times though we seem to hit a limit on our ability to even get good approximations. We now know that we cannot achieve better approximations on many of these problems unless **P** = **NP** and we could solve these problems exactly. The techniques to show these negative results came out of a new model of proof system originally developed for cryptography and to classify group theoretic algorithmic problems.

As mentioned in Section 4.3 we don't expect to have short traditional proofs of tautologies. But consider an “interactive proof” model where a prover Peggy tries to convince a verifier Victor that a formula ϕ is a tautology. Victor can ask Peggy randomly generated questions and need only be convinced with high confidence. Quite surprisingly these proof systems have been shown to exist not only for tautologies but for any problem computable in a reasonable amount of memory.

A variation known as a “probabilistically checkable proof system” (PCPs), where Peggy writes down an encoded proof and Victor can make randomized queries to the bits of the proof, has applications for approximations. The “PCP Theorem” which optimizes parameters which in its strong form shows that every language in **NP** has a PCP where Victor uses a tiny number of random coins and queries only three bits of the proof.

One can use this PCP theorem to show the limitations of approximation for a large number of optimization questions. For example one cannot approximate the largest clique in a group of n people by more than a multiplicative ratio of nearly \sqrt{n} unless **P** = **NP**. See the recent CACM survey by Madhu Sudan for more details and references on PCPs [37].

One can do even better assuming a “Unique Games Conjecture,” that there exists PCPs for **NP** problems with some stronger properties. Consider the MAX-CUT problem of dividing people discussed in Section 5.3. If the unique games conjecture holds one cannot do better than the .878567 factor given by the Goemans-Williamson approximation algorithm [26]. Recent work shows how to get a provably best approximation for essentially any constrained problem assuming this conjecture [31].

7. USING HARDNESS

In Section 3 we saw the nice world that arises when we assume **P** = **NP**. But we expect **P** \neq **NP** to hold in very strong ways. We can use strong hardness assumptions as a positive tool, particularly to create cryptographic protocols and to reduce or even eliminate the need of random bits in probabilistic algorithms.

7.1 Cryptography

We take it for granted these days, the little key or lock on our web page that tells us that someone listening to the network won't get the credit card number I just sent to an on-line store or the password to the bank that controls my money. But public-key cryptography, the ability to send secure messages between two parties that have never privately exchanged keys, is a relatively new development based on hardness assumptions of computational problems.

If **P** = **NP** then public-key cryptography is impossible. As-

suming $\mathbf{P} \neq \mathbf{NP}$ is not enough to get public-key protocols, instead we need strong average-case assumptions about the difficulty of factoring or related problems.

We can do much more than just public-key cryptography using hard problems. Suppose Alice’s husband Bob is working on a Sudoku puzzle and Alice claims she has a solution to the puzzle (solving a $n^2 \times n^2$ -Sudoku puzzle is \mathbf{NP} -complete). Can Alice convince Bob that she knows a solution without revealing any piece of it?

Turns out Alice can using a “zero-knowledge proof,” an interactive proof with the additional feature that the verifier learns nothing other than some property holds, like a Sudoku puzzle having a solution. Every \mathbf{NP} search problem has a zero-knowledge proof under the appropriate hardness assumptions.

On-line poker is generally played through some “trusted” website, usually somewhere in the Caribbean. Can we play poker over the Internet without a trusted server? Using the right cryptographic assumptions, not only poker but any protocol that uses a trusted party can be replaced by one that uses no trusted party and the players can’t cheat or learn anything new beyond what they could do with the trusted party.

See the survey of Goldreich [19] for details on the above.

7.2 Eliminating Randomness

In the 70’s we saw a new type of algorithm, one that used random bits to aid in finding a solution to a problem. Most notably we had probabilistic algorithms [36] for determining whether a number is prime, an important routine needed for modern cryptography. In 2004 we discovered we don’t need randomness at all to efficiently determine if a number is prime [3]. Does randomness help us at all in finding solutions to \mathbf{NP} problems?

Truly independent and uniform random bits are either very difficult or impossible to produce (depending on your beliefs about quantum mechanics). Computer algorithms instead use pseudorandom generators to generate a sequence of bits from some given seed. The generators typically found on our computers usually work well but occasionally give incorrect results both in theory and in practice.

We can create theoretically better pseudorandom generators in two different ways, one based on the strong hardness assumptions of cryptography and the other based on worst-case complexity assumptions. I will focus on this second approach.

We need to assume a bit more than $\mathbf{P} \neq \mathbf{NP}$, roughly that \mathbf{NP} -complete problems cannot be solved by smaller than expected AND-OR-NOT circuits. A long series of papers showed that, under this assumption, any problem with an efficient probabilistic algorithm also has an efficient algorithm that uses a pseudorandom generator with a very short seed, a surprising connection between hard languages and pseudorandomness (see [24]). The seed is so short we can try all possible seeds efficiently and avoid the need for randomness altogether.

Thus complexity theorists generally believe having randomness does not help in solving \mathbf{NP} search problems and that \mathbf{NP} -complete problems do not have efficient solutions, either with or without using truly random bits.

While randomness doesn’t seem necessary for solving search problems, the unpredictability of random bits plays a critical role in cryptography and interactive proof systems and likely cannot be avoided in these scenarios.

8. COULD QUANTUM COMPUTERS SOLVE \mathbf{NP} -COMPLETE PROBLEMS?

While we have randomized and non-randomized efficient algorithms for determining whether a number is prime, these algorithms usually don’t give us the factors of a composite number. Much of modern cryptography relies on the fact that factoring or similar problems do not have efficient algorithms.

In the mid-90’s, Peter Shor [35] showed how to factor numbers using a hypothetical quantum computer. He also gave a similar quantum algorithm to solve the discrete logarithm problem. The hardness of discrete logarithm on classical computers is also used as a basis for many cryptographic protocols. Nevertheless we don’t expect that factoring or finding discrete logarithms are \mathbf{NP} -complete. While we don’t think we have efficient algorithms to solve factoring or discrete logarithm, we also don’t believe we can reduce \mathbf{NP} -complete problems like Clique to the factoring or discrete logarithm problems.

So could quantum computers one day solve \mathbf{NP} -complete problems? Unlikely.

I’m not a physicist so I won’t address the problem as to whether these machines can actually be built at a large enough scale to solve factoring problems larger than we can with current technology (about 200 digits). After billions of dollars of funding of quantum computing research we still have a long way to go.

Even if we could build these machines, Shor’s algorithm relies heavily on the algebraic structures of numbers that we don’t see in the known \mathbf{NP} -complete problems. We know that his algorithm cannot be applied to generic “black-box” search problems so any algorithm would have to use some special structure of \mathbf{NP} -complete problems that we don’t know about. We have used some algebraic structure of \mathbf{NP} -complete problems for interactive and zero-knowledge proofs but quantum algorithms would seem to require much more.

Lov Grover [20] did find a quantum algorithm that works on general \mathbf{NP} problems but that algorithm only achieves a quadratic speed-up and we have evidence that those techniques will not go further.

Meanwhile quantum cryptography, using quantum mechanics to achieve some cryptographic protocols without hardness assumptions, has had some success both in theory and in practice.

9. A NEW HOPE?

Ketan Mulmuley and Milind Sohoni have presented an approach to the **P** vs. **NP** problem through algebraic geometry, dubbed Geometric Complexity Theory, or GCT [30]. This approach seems to avoid the difficulties mentioned in Section 4 but requires deep mathematics that could require many years or decades to carry through.

In essence, they define a family of high-dimension polygons P_n based on group representations on certain algebraic varieties. Roughly speaking, for each n , if P_n contains an integral point, then any circuit family for the Hamiltonian path problem must have size at least $n^{\log n}$ on inputs of size n , which implies **P** \neq **NP**. Thus, to show that **P** \neq **NP** it suffices to show that P_n contains an integral point for all n .

Although all that is necessary is to show that P_n contains an integral point for all n , Mulmuley and Sohoni argue that this direct approach would be difficult and instead suggests first showing that the integer programming problem for the family P_n is in **P**. Under this approach, there are three significant steps remaining:

1. Prove that the LP relaxation solves the integer programming problem for P_n in polynomial time,
2. find an efficient, simple combinatorial algorithm for the integer programming problem for P_n , and
3. prove that this simple algorithm always answers “yes.”

Since the polygons P_n are algebro-geometric in nature, solving (1) is thought to require algebraic geometry, representation theory, and the theory of quantum groups. Mulmuley and Sohoni have given reasonable algebro-geometric conditions that imply (1). These conditions have classical analogues that are known to hold, based on the Riemann Hypothesis over finite fields (a theorem proved by André Weil in the 1960s). Mulmuley and Sohoni suggest that an analogous Riemann Hypothesis-like statement is required here (though *not* the classical Riemann Hypothesis).

Although step (1) is difficult, Mulmuley and Sohoni have provided definite conjectures based on reasonable mathematical analogies that would solve (1). In contrast, the path to completing steps (2) and (3) is less clear. Despite these remaining hurdles, even solving the conjectures involved in (1) could provide some insight to the **P** versus **NP** problem.

Mulmuley and Sohoni have reduced a question about the *non*-existence of polynomial-time algorithms for all **NP**-complete problems to a question about the *existence* of a polynomial-time algorithm (with certain properties) for a specific problem. This should give us some hope, even in the face of problems (1)-(3).

Nevertheless, Mulmuley himself believes it will take about 100 years to carry out this program, if it works at all.

10. CONCLUSION

This survey focused on the **P** versus **NP** problem, its importance, our attempts to prove **P** \neq **NP** and the approaches we use to deal with the **NP**-complete problems that nature

and society throws at us. Much of the work mentioned required a long series of mathematically difficult research papers that we could not hope to adequately cover in this short article. Also the field of computational complexity goes well beyond just the **P** versus **NP** problem that we haven't discussed here. In Section 11 we give a number of references for those interested in a deeper understanding of the **P** versus **NP** problem and computational complexity.

The **P** versus **NP** problem has gone from an interesting problem related to logic to perhaps the most fundamental and important mathematical question of our time, whose importance only grows as computers become more powerful and widespread. The question has even hit popular culture appearing in television shows such as *The Simpsons* and *Numb3rs*. Yet many only know of the basic principles of **P** versus **NP** and I hope this survey has given you a small feeling of the depth of research inspired by this mathematical problem.

Proving **P** \neq **NP** would not be the end of the story, it would just show that **NP**-complete problems don't have efficient algorithms for all inputs but many questions might remain. Cryptography for example would require that a problem like factoring (not believed to be **NP**-complete) is hard for randomly drawn composite numbers.

Proving **P** \neq **NP** might not be the start of the story either. Weaker separations remain perplexingly difficult, for example showing that Boolean-formula Satisfiability cannot be solved in near-linear time or showing that some problem using a certain amount of memory cannot be solved using roughly the same amount of time.

None of us truly understand the **P** versus **NP** problem, we have only begun to peel the layers around this increasingly complex question. Perhaps we will see a resolution of the **P** versus **NP** problem in the near future but I almost hope not. The **P** versus **NP** problem continues to inspire and boggle the mind and continued exploration of this problem will lead us to yet even new complexities in that truly mysterious process we call computation.

11. FURTHER READING

Recommendations for a more in-depth look at the **P** versus **NP** problem and the other topics discussed in this article:

- Steve Homer and myself have written a detailed historical view of computational complexity [15].
- The 1979 book of Garey and Johnson still gives the best overview of the **P** versus **NP** problem with an incredibly useful list of **NP**-complete problems [17].
- Scott Aaronson looks at the unlikely possibility that the **P** versus **NP** problem is formally independent [2].
- Russell Impagliazzo gives a wonderful description of five possible worlds of complexity [23].
- Sanjeev Arora and Boaz Barak have a new computational complexity textbook with an emphasis on recent research directions [6].

- The Foundations and Trends in Theoretical Computer Science journal and the Computational Complexity columns of the Bulletin of the European Association of Theoretical Computer Science and SIGACT News have many wonderful surveys on various topics in theory including those mentioned in this article.
- Read the weblog Computational Complexity and you will be among the first to know on any updates of the status of the **P** versus **NP** problem [14].

Acknowledgments

Thanks to Rahul Santhanam for many useful discussions and comments. Josh Grochow wrote an early draft of Section 9. The anonymous referees provided critical advice. Some of the material in this article has appeared in my earlier surveys and my blog [14].

12. REFERENCES

- [1] The international SAT competitions web page. <http://www.satcompetition.org>.
- [2] S. Aaronson. Is P versus NP formally independent? *Bulletin of the European Association for Theoretical Computer Science*, 81, Oct. 2003.
- [3] M. Agrawal, N. Kayal, and N. Saxena. PRIMES is in P. *Annals of Mathematics*, 160(2):781–793, 2004.
- [4] D. Applegate, R. Bixby, V. Chvátal, and W. Cook. On the solution of traveling salesman problems. *DOCUMENTA MATHEMATICA*, Extra Volume ICM III:645–656, 1998.
- [5] S. Arora. Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems. *Journal of the ACM*, 45(5):753–782, Sept. 1998.
- [6] S. Arora and B. Barak. *Complexity Theory: A Modern Approach*. Cambridge University Press, Cambridge, 2009.
- [7] T. Baker, J. Gill, and R. Solovay. Relativizations of the P = NP question. *SIAM Journal on Computing*, 4(4):431–442, 1975.
- [8] G. Cantor. Ueber eine Eigenschaft des Inbegriffes aller reellen algebraischen Zahlen. *Crelle's Journal*, 77:258–262, 1874.
- [9] B. Cipra. This ising model is NP-complete. *SIAM News*, 33(6), July/August 2000.
- [10] V. Conitzer and T. Sandholm. New complexity results about Nash equilibria. *Games and Economic Behavior*, 63(2):621–641, July 2008.
- [11] S. Cook. The complexity of theorem-proving procedures. In *Proceedings of the 3rd ACM Symposium on the Theory of Computing*, pages 151–158. ACM, New York, 1971.
- [12] R. Downey and M. Fellows. *Parameterized Complexity*. Springer, 1999.
- [13] J. Edmonds. Paths, trees and flowers. *Canadian Journal of Mathematics*, 17:449–467, 1965.
- [14] L. Fortnow and W. Gasarch. Computational complexity. <http://weblog.fortnow.com>.
- [15] L. Fortnow and S. Homer. A short history of computational complexity. *Bulletin of the European Association for Theoretical Computer Science*, 80, June 2003. Computational Complexity Column.
- [16] M. Furst, J. Saxe, and M. Sipser. Parity, circuits and the polynomial-time hierarchy. *Mathematical Systems Theory*, 17:13–27, 1984.
- [17] M. Garey and D. Johnson. *Computers and Intractability. A Guide to the theory of NP-completeness*. W. H. Freeman and Company, New York, 1979.
- [18] M. Goemans and D. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM*, 42(6):1115–1145, 1995.
- [19] O. Goldreich. Foundations of cryptography—a primer. *Foundations and Trends in Theoretical Computer Science*, 1(1):1–116, 2005.
- [20] L. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the 28th ACM Symposium on the Theory of Computing*, pages 212–219. ACM, New York, 1996.
- [21] D. Gusfield. *Algorithms on Strings, Trees and Sequences: Computer Science and Computational Biology*. Cambridge University Press, 1997.
- [22] A. Haken. The intractability of resolution. *Theoretical Computer Science*, 39:297–305, 1985.
- [23] R. Impagliazzo. A personal view of average-case complexity theory. In *Proceedings of the 10th Annual Conference on Structure in Complexity Theory*, pages 134–147. IEEE Computer Society Press, 1995.
- [24] R. Impagliazzo and A. Wigderson. P = BPP if E requires exponential circuits: Derandomizing the XOR lemma. In *Proceedings of the 29th ACM Symposium on the Theory of Computing*, pages 220–229. ACM, New York, 1997.
- [25] R. Karp. Reducibility among combinatorial problems. In R. Miller and J. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103. Plenum Press, 1972.
- [26] S. Khot, G. Kindler, E. Mossel, and R. O’Donnell. Optimal inapproximability results for MAX-CUT and other 2-variable CSPs? *SIAM Journal on Computing*, 37(1):319–357, 2007.
- [27] R. Lathrop. The protein threading problem with sequence amino acid interaction preferences is NP-complete. *Protein Engineering*, 7(9):1059–1068, 1994.
- [28] L. Levin. Universal’nyie perebornyie zadachi (Universal search problems: in Russian). *Problemy Peredachi Informatsii*, 9(3):265–266, 1973. Corrected English translation in [38].
- [29] L. Levin. Average case complete problems. *SIAM Journal on Computing*, 15:285–286, 1986.
- [30] K. Mulmuley and M. Sohoni. Geometric complexity theory I: An approach to the P vs. NP and related problems. *SIAM Journal on Computing*, 31(2):496–526, 2001.
- [31] P. Raghavendra. Optimal algorithms and inapproximability results for every csp? In *Proceedings of the 40th ACM Symposium on the Theory of Computing*, pages 245–254, New York, NY, USA, 2008. ACM.
- [32] A. Razborov. Lower bounds on the monotone

- complexity of some boolean functions. *Soviet Mathematics–Doklady*, 31:485–493, 1985.
- [33] A. Razborov. On the method of approximations. In *Proceedings of the 21st ACM Symposium on the Theory of Computing*, pages 167–176. ACM, New York, 1989.
- [34] A. Razborov and S. Rudich. Natural proofs. *Journal of Computer and System Sciences*, 55(1):24–35, Aug. 1997.
- [35] P. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, 1997.
- [36] R. Solovay and V. Strassen. A fast Monte-Carlo test for primality. *SIAM Journal on Computing*, 6:84–85, 1977. See also erratum 7:118, 1978.
- [37] M. Sudan. Probabilistically checkable proofs. *Communications of the ACM*, 52(3):76–84, Mar. 2009.
- [38] R. Trakhtenbrot. A survey of Russian approaches to *Perebor* (brute-force search) algorithms. *Annals of the History of Computing*, 6(4):384–400, 1984.
- [39] A. Turing. On computable numbers, with an application to the Entscheidungs problem. *Proceedings of the London Mathematical Society*, 42:230–265, 1936.
- [40] D. van Melkebeek. A survey of lower bounds for satisfiability and related problems. *Foundations and Trends in Theoretical Computer Science*, 2(197-303), 2007.